DOCUMENT RESUME

ED 108 511 FL 006 982

AUTHOR Rosenberg, Richard S.
TITLE Artificial Intelligence and Linguistics: A Brief History of a One-Way Relationship.
PUB DATE Feb 75
NOTE 15p.; Paper presented at the Annual conference of the Berkeley Linguistics Society (1st, Berkeley, California, February 1975)

EDRS PRICE MF-$0.76 HC-$1.58 PLUS POSTAGE
DESCRIPTORS *Artificial Intelligence; *Computational Linguistics; *Computer Programs; *Computer Science; Information Theory; Linguistics; *Linguistic Theory; Machine Translation; Models; Programing Languages; Semantics; Systems Development

ABSTRACT

For the past 15 years there has been a serious interest in the processing of natural language (English) by researchers in Artificial Intelligence (A.I.). This processing has included machine translation, question-answering systems, man-machine dialogue, and speech understanding. This interest has engendered an awareness of and a concern with the ongoing activity in contemporary linguistics. Therefore, it may be of interest to linguistis to discover what has seemed important for A.I. and how it has been adapted and used. Thus a brief history of the relation (almost always one-way) between A.I. and linguistics is presented. Some of the works in A.I. surveyed range from those of the early sixties, such as Lindsay's SAD SAM, Green et al BASEBALL, and Bobrow's STUDENT, to more recent efforts, including Wood's transition network grammars, Winograd's SHRDLU, and Schank's conceptual dependency models. In one way or another, these computer programs and others depend on the work of Chomsky, both "Syntactic Structures" and "Aspects," Halliday's systemic grammar, and some of the ideas of generative semantics as developed by G. Lakoff, McCawley, and Fillmore. (Author/KM)

****
* Documents acquired by ERIC include many informal unpublished *
* materials not available from other sources. ERIC makes every effort *
* to obtain the best copy available. nevertheless, items of marginal *
* reproducibility are often encountered and this affects the quality *
* of the microfiche and hardcopy reproductions ERIC makes available *
* via the ERIC Document Reproduction Service (EDRS). EDRS is not *
* responsible for the quality of the original document. Reproductions *
* supplied by EDRS are the best that can be made from the original. *
****

ED108517

Permission to reproduce this copyrighted material has been granted to Berkeley Linguistics Society and to the organization operating under agreements with the National Institute of Education further reproduction outside the ERIC system requires permission of the copyright owner.

Proceedings of
First Annual Conference of the
Berkeley Linguistics Society
Feb. 1975, Berkeley, CA.

U.S. DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION
THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

ARTIFICIAL INTELLIGENCE AND LINGUISTICS: A Brief History of a One-Way Relationship

Richard S. Rosenberg
The University of British Columbia

1. Introduction

For the past fifteen years[1] there has been serious interest in the processing of natural language (English) by researchers in Artificial Intelligence (A.I.). This processing has included machine translation, question-answering (Q-A) systems, man-machine dialogue, and speech understanding. This keen interest has engendered an awareness of and a concern with ongoing activity in contemporary linguistics. Therefore, it may be of interest to linguists to discover what has seemed important for A.I. and how it has been adapted and used. Furthermore, I would suggest that current activity in A.I. should be of interest in its own right to linguists concerned with the whole range of problems involved in human language understanding.

An enduring definition of the goal of A.I. research is "to construct computer programs which exhibit behavior that we call 'intelligent behavior' when we observe it in human beings." (Feigenbaum and Feldman 1963, p.3) Winograd (1971), in discussing the goals of his research--which include the desire to have a usable language-understanding system and to gain a better understanding of what language is and how it is put together--states:

> More generally, we want to understand what intelligence is and how it can be put into computers. Language is one of the most complex and unique of human activities, and understanding its structure may lead to a better theory of how our minds work.

These remarks suggest that there should be closer cooperation among A.I. researchers and linguists, although this clearly has not been the case, as Walker (1973) notes:

> Linguists for the most part have not accepted the computer or even computation as an essential methodological component of their field. Moreover, many linguists have denied not only the relevance of the results of computational linguistics research for linguistics, but, more importantly, the possible relevance of such results.

FL00 6982

I suspect this situation may derive from positions which can best be expressed by two hypothetical researchers:

A.I. Researcher: Linguists are too preoccupied with very small aspects of the total language problem. They need to look at the larger picture and they also need the realities of a computer program to constrain their imaginations in order to produce a more precise formulation of their theories.

Linguist: So far, A.I. researchers have trivialized the complexities of language understanding. They write large programs dealing with narrow domains involving relatively simple grammatical constructions. They just don't appreciate that we know so very little that a lot of

basic research is of the highest priority.

From the outset, the A.I. researcher has been compelled to deal with problems of syntax, semantics, context, pragmatics, and representation of knowledge for the design of even the simplest system. And although the treatment has often appeared superficial, recognition of the simplifying assumptions made and the difficulties of the problem has always been made. It is my task here to amplify this point of view.

The following sections will consist of a brief discussion of machine translation, a more extended treatment of Q-A systems, and finally a look at current research in speech understanding.

2. Machine Translation

Machine translation as a modern enterprise began in the late 1940's, and as of the present date it can be said that fully automatic high quality translation (FAHQT) has not been achieved nor are the prospects hopeful. The major effort in machine translation was a precursor to A.I. research, and its failures have been used to cast doubt upon the whole A.I. endeavour (see Dreyfus 1972). As early as 1959, Bar Hillel (1960), one of the earliest workers in this field and later its severest critic, argued that FAHQT was not a reasonable goal and was even an impossible one. Since much of the work up to that time had made use of rather simple grammatical notions together with large dictionaries and was basically word for word, his position was not altogether unreasonable.

But Bar Hillel did comment on the optimism aroused by the then recent achievements of Noam Chomsky (1957) and the hopes that his transformational model would lead to success where previous syntactical analysis had not. However, much of the work in mechanical translation continued to be based on the notion that words are the units of meaning and to this idea were added "notions of thesaurus classes of words, statistical associations, probabilities and superficial syntactic structures" (Simmons 1970). The final blow came with the issue of the ALPAC (1966) report, which signaled the effective end of large-scale support for machine translation.

One legacy of machine translation was a growing appreciation of the complexities of language and an awareness of the necessity to integrate various kinds of knowledge if language was to be dealt with in any meaningful manner. Recently there has been a revival of interest (see Wilks 1973a, 1973b). It should be mentioned that machine-aided translation involving either pre-editing or post-editing is an alternative but not a particularly attractive one to researchers in A.I.

3. Question-Answering Systems

For many of the early systems, the decision was made to avoid the necessity of dealing with the full complexity of natural language by the use of specially designed formats for both the input and the representation of knowledge within the system. Because the following treatment will of necessity be brief, the reader would be well advised to consult such surveys as Bobrow (1963), Simmons (1965,1970), Kuno (1967), Walker (1973a), and more detailed sources such as Feigenbaum

and Feldman (1963), Minsky (1968), Winograd (1972), Schank and Colby (1973), and Rustin (1973).

BASEBALL (Green et al 1961)

This computer program answers questions posed in ordinary English about wins and losses in American league baseball games. Typical questions are:

Where did the Red Sox play on July 7?
What teams won 10 games in June?
How many games did the Yankees play in May?

For our present purposes it is important to note that the linguistic routines involve a dictionary look-up for parts of speech and definitions, a syntactic analysis based on the work of Zelig Harris (1960) which is usually successful, in this case because of the restricted format for questions, and a semantic routine which uses the dictionary meanings and the results of the syntactic analysis to produce a specification list with which the fixed data structure can be searched. This program works well within its very limited domain, especially with such restrictions on the input sentences as single clauses, prohibition of "and", "or", and "not" as well as constructions involving relations such as "most" and "highest". But basically the level of language understanding is quite limited.

SAD SAM Sentence Appraiser and Diagrammer and Semantic Analyzing Machine (Lindsay 1963)

The basic semantic information of interest is family relationships, e.g. father, mother, brother, etc. There are eight such relationships in Basic English, a 1700-word subset of English which provides the acceptable vocabulary for the program. The task of the program is to construct family trees by extracting the kinship relationships contained in the input sentences. Sentences are analyzed by means of a context-free grammar, and the parse tree is used to extract the contained kinship relation. Since the semantics are limited to kinship terms, everything else in the parse is ignored, so that a sentence like

John visited his sister Mary in Chicago during the summer of 1967.

will result in a structure being built which represents only the information that John and Mary are siblings.

This program foreshadows the dominant linear processing paradigm of the 1960's: First parse to produce a representation of the input sentence, usually a tree; then apply "semantic" routines to produce a query language statement; finally, execute to retrieve information from the data base (see Kuno 1967). The use of grammars expressed as systems of rules is entirely due to Chomsky, and his influence is very pervasive in this period.

STUDENT (Bobrow 1964)

STUDENT is a program designed to accept an English language statement of a high school algebra problem, to convert it into a set of simultaneous linear equations, and finally to solve this set (if possible) and produce an answer. A typical problem is

> If the number of customers Tom gets is twice the square of 20 per cent of the number of advertisements he runs, and the number of advertisements he runs is 45, what is the number of customers Tom gets?

Bobrow makes use of a notion of kernel sentences and transformations which he claims are different from Chomsky's; however, his idea is to assume an underlying structure which must be uncovered by a small set of sentence forms. In this restricted domain, meaning is expressible as a set of equations of the form P1=P2, where P1 and P2 may be strings of uninterpreted symbols. The program contains a number of procedures necessary to carry out simple pronoun reference, to match strings which are formally different, to relate the subject of one sentence to its immediate predecessor in the text, etc. Because of the restricted domain, a fairly small set of prestored sentence forms is sufficient to deal with a large number of apparently different sentences. A somewhat extended version of this notion is the basis for programs called by various names such as Doctor and ELIZA (Weizenbaum 1966).

Logic-Based Systems

In 1958, John McCarthy presented a paper titled "Programs with Common Sense" in which he proposed a program called Advice Taker, which would solve problems by manipulating sentences in formal languages. This motivated research in at least three directions: the development of adequate formal representations for natural language sentences, the investigation of problem solving in formal systems, and the study of methods for translating from natural language into formal languages.

Raphael (1964) developed a program called SIR (Semantic Information Retrieval) with the goal being to derive answers to simple questions expressed in natural language by the use of deductive procedures. As such, the class of input sentences was quite restricted, being limited to set relations, ownership, part-whole, number, and position. Employing a large number of special purpose routines and operating on a symbolic representation of the input information, the program was able to perform rather simple deductions.

In the mid-1960's, a powerful proof procedure for predicate calculus was developed by J.A. Robinson (1965). Thus, if problems could be formalized in predicate calculus, this procedure, called resolution, offered the possibility of an effective theorem prover as a problem solver. It should also be mentioned that during this period as well there was a considerable research effort devoted to robot systems which it was hoped could be commanded in English to carry out non-trivial tasks requiring some form of deduction. A new

approach thus emerged which required the translation of English into predicate calculus, the application of a theorem prover to the resulting formula, and, if successful, the return of a solution. For example, the input

Is there a green box under the table?

would be translated into something like

$$(\exists x,y)[\mathrm{Box}(x) \wedge \mathrm{Green}(x) \wedge \mathrm{Table}(y) \wedge \mathrm{Under}(x,y)]$$

The theorem prover must now determine if there are objects x and y which satisfy the above formula and if there are, the answer to the question will be yes with the specific objects named. Of course, this approach requires all information about the domain to be stored as predicate calculus formulae.

A problem requiring deduction for solution, for example, a version of the famous 'monkey and bananas' problem, will return the sequence of operations necessary for a solution (go to the box, push the box under the bananas, climb the box, get the bananas). The most important work done in this period was Green and Raphael (1968). Other interesting efforts using predicate calculus are Sandewall (1971) and Coles (1972). However, this overall approach has lost much favour, because for large data bases, the theorem prover, employing various refinements of the resolution principle, has turned out to be rather inefficient.

It is interesting that there has been a simultaneous revival of interest in the formal modeling of natural language in philosophy and linguistics as well as A.I. In linguistics the names of George Lakoff (1970) and James McCawley (1968), 1969, 1972) are particularly prominent. It must be noted however that on the whole their work has had little influence in A.I. Almost the only reference to it is to be found in Wilks (1972, 1973b). In the former he expresses strong disagreement with Lakoff's approach. In fact, much of the effort of the generative semanticists has gone virtually unnoticed in the A.I. community. Chomsky (1965) has continued to exert considerable influence even though his theories have come under strong and continual attack since 1966. This in itself is noteworthy, for with his emphasis on syntax to the exclusion of semantics, pragmatics, and much else, his point of view is in direct opposition to the requirements of Q-A systems. Nevertheless, there is a certain attraction to Chomsky's framework for syntax which lends itself to computer implementation. It is not surprising that the unsettled state of generative semantics should discourage the interest of researchers in A.I.

Augmented Transition Network Grammars

The names of W. Woods (1970, 1973) and of Woods and Kaplan (1971) are most prominently associated with this approach to natural language processing. Woods (1968) proposed the use of procedural semantics for a Q-A system. Assuming that the program had available a deep structure parse or phrase marker (in the Chomsky sense) of the

input question, he developed a system which attempted to match parts of the tree with a collection of prestored sub-trees. Each successful match would result in the incorporation of a piece of LISP code into a growing program. The execution of the completed program results in an answer to the original question. An example of a question and its semantic interpretation in terms of predicates and functions is

> What is the departure time from Boston of every American Airlines flight that goes from Boston to Chicago?
> (FOR EVERY X2/ FLIGHT: EQUAL (OWNER (X2), AMERICAN AIRLINES) AND CONNECT (X2,BOSTON,CHICAGO); (FOR THE X1/DTIME (X2,BOSTON): LIST (X1)))

The augmented transition network (ATN) grammar is a formulation of a transformational grammar which produces the (Chomsky-like) deep structure parse for input to the semantic processor. More detail is found in Woods (1970, 1973), but for the present it can be said that an ATN grammar is a collection of finite state graphs which allow recursive calls on the arcs and which has provision for storing partial parses. Some of its virtues are the clarity of the representation, the speed of processing, and the way in which regularities of language are captured by having a single piece of network serve multiple uses. The Woods system probably carries the linear paradigm for natural language processing as far as it can go.

Most linguists of a generative semantics persuasion would probably dismiss this whole effort as misguided and hopelessly inadequate as a model for language processing. But it does provide a fairly powerful system for certain practical situations. A prototype of Wood's system was designed as a Q-A fact retrieval program to answer a geologist's questions about the reported chemical analyses of lunar rock samples brought back from Apollo missions. With a data base of approximately 13,000 individual chemical analysis measurements, the program performed remarkably well when run during a Lunar Science Conference in 1971. (See Woods 1973) Sample questions are

> Do any breccias contain Aluminum?
> What are they?
> In how many breccias is the average concentration of Aluminum greater than 13 per cent?

It is also clear that the construction of such large systems does reveal many factors involved in the task of natural language processing and reveals them in a particularly immediate fashion.

SHRDLU (Winograd 1971, 1972, 1973)

This unlikely string of symbols is the name of Terry Winograd's very important program for natural language understanding. A major departure is made by Winograd from the by now traditional generative grammar approach, to a view of language enunciated by M.A.K. Halliday (1970). There is a strong similarity between Halliday's position

and that of Fillmore (1968) to whom we shall return subsequently. Halliday's theory views meaning as having a central place in the way language is structured, and he proposes a number of "system networks" which describe how different features of a sentence interact and depend on one another. This is reflected in the following (Winograd 1971):

> If we really want computers to understand us, we need to give them the ability to use more knowledge. In addition to a grammar of the language, they need to have all sorts of knowledge about the subject they are discussing, and they have to use all sorts of reasoning to combine facts in the right way to understand and respond to it. The process of understanding a sentence has to combine grammar, semantics, and reasoning in a very intimate way, calling on each part to help with the others.

It is this last sentence which expresses the crucial aspect of Winograd's contribution to natural language processing. Sharp distinctions between the various phases of processing--syntax, semantics, inference, context--are done away with, resulting in a system that is difficult to describe but powerful in its operation.

SHRDLU's domain of interest is a simulation of a hand-eye system resembling those that have been built at Stanford and M.I.T. For the simulation, we can visualize a table with blocks of various sizes, shapes, and colours, a box, and a mechanical arm able to move one block at a time. Thus the system can respond to commands ("Pick up a red block"), questions ("What does the box contain?"), and declaratives ("A 'steeple' is a stack which contains two green cubes and a pyramid."). The program requires the ability to perform inference when dealing with a command such as

> Find a block which is taller than the one you are holding and put it into the box.

In addition, it must determine what is meant by "one" and "it". It can also deal with logical connectives, both in the grammar and semantics

> Will you please stack up both of the red blocks and either a green cube or a pyramid.

as well as a wide range of grammatical constructions including passives, coordinates, and comparatives.

The inference powers of the system are lodged in programs written in the language MICROPLANNER which is a procedural language for doing various kinds of problem solving. What is interesting in Winograd's system is that inference plays an integral role in processing the natural language as well as in the associated problem solving. The role of the semantics component is to translate from grammatical structures into MICROPLANNER code. This process, however, is not done at the end of a grammatical phase but is carried out in tandem with

the grammatical analysis.

Since the program saves the ongoing dialogue, it is possible to question it about past events, and the program itself can use the dialogue to resolve questions of reference and context. In terms of its scope both with respect to linguistic problems and problems of representation of knowledge, linguists would be well-advised to become familiar with Winograd's SHRDLU.

4. Other Systems

I want to mention two projects which do not fit under the category Q-A systems. These are Simmons' (1973) work on semantic networks and Schank's (1972, 1973a, 1973b) conceptual analysis models.

Simmons' model is based very strongly on Fillmore's (1968) theory of deep case structures with further developments by Celce-Murcia (1972) to provide a form for representing knowledge together with an ATN grammar which is used to transform the input English into these semantic structures. Fillmore has probably been the most influential linguist in A.I. over the past few years, with several systems based on his theories. However, some of his recent writings (Fillmore 1971) which raise questions about case grammars do not seem to have become well known in A.I. We might also mention Fillmore's influence in cognitive psychology, especially in the memory models of Norman (1973), Rumelhart et al (1972), and Kintsch (1972). The interaction of A.I. researchers and cognitive psychologists in such shared interests as memory, knowledge, semantics in particular and language understanding and production in general, is an encouraging feature of the past few years[2]. The addition of linguists to this enterprise would certainly be a welcome event.

Simmons has also been noteworthy for his concern with sentence generation and has developed an ATN for generating sentences (in response to questions) using information from the appropriate part of the semantic nets. An input question is processed by an ATN grammar to produce a deep case structure parse. Search routines attempt to match this structure with knowledge also prestored in semantic networks of case structures. When matches are found, they are input to the sentence generator component which produces an appropriate response.

Schank and his students have been trying to produce a model of human language understanding to serve as the basis for a computer program and have indeed produced a prototype system. The program is a kind of language understander which upon being presented with an input will produce a series of paraphrases and will also draw some simple inferences. The major stress here is on underlying meaning or conceptualizations, with a forcefully expressed lack of interest in syntactic constructions. Grammar is seen as a minor aid in determining meaning, although the system does have a conceptual parser as a first phase. The aim is to produce a representation pregnant with associations and potentialities for inference.

Schank uses the term ACT to represent the underlying actions implied by the sentence surface structure. There is certainly a similarity between his notion of a conceptual dependency network

centered around ACT's and Fillmore's case frames, although Schank argues that Fillmore's work is essentially syntactic. As an example of the underlying complexity of an apparently simple sentence, consider Schank's representation of

I walked to the cafeteria.

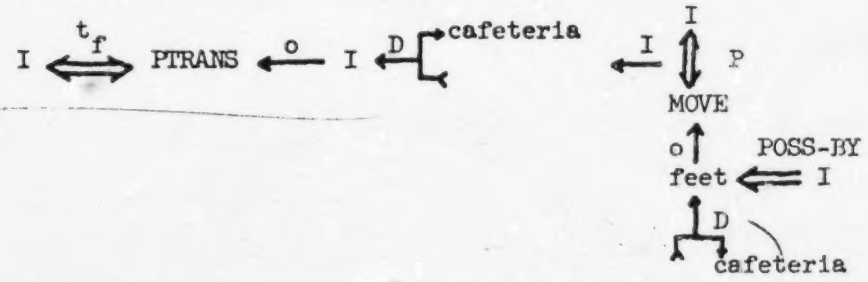

I will not attempt to explain the notation but basically the sense is that "I completed a change in my position to the cafeteria by means of (instrument I) having moved my feet to the cafeteria." Thus a valid inference which could be drawn here is that the object (I) is located at the direction (cafeteria). PTRANS and MOVE are two of fourteen primitive ACT's which Schank claims are necessary to represent all the actions underlying natural language. His system incorporates large numbers of rules which operate on the input sentence using a large and complex lexicon to produce conceptual structures. Schank will settle for nothing less than a model which explains how people understand natural language.

5. Speech Understanding Systems

At present there are three large-scale research efforts to achieve working speech systems by 1978: Reddy et al (1973), Woods and Makhoul (1973), and Walker (1973b, 1974). The basic specifications for these projects, all of which are supported by the Advanced Research Projects Agency, were drawn up in a report issued by Newell et al (1973). Some of these are that the system should accept continuous speech from many speakers of general American dialect, in a quiet room, over a good quality microphone, with a vocabulary of about 1000 words and a highly artificial syntax. The domain of discourse should be limited; the system, slightly tuned for each speaker, should operate in real time with less than 10% semantic error.

The development of such systems is a great challenge, since all the complexities of language understanding must be dealt with in relating acoustic signals to meaning. As such, researchers on these projects have, of necessity, to draw from the extensive phonological literature and integrate work in this area with current ideas from Q-A systems already in operation. But now the linear paradigm is totally inadequate because there is no way to convert the acoustic signal to a sequence of words without a whole range of syntactic, semantic, context and problem domain knowledge. It would seem that any single aspect of the language processing task requires informa-

tion and help from all other aspects.

Winograd (1974) has described two approaches to system organization for dealing with the apparent necessity of an all-at-once approach. The first he calls heterarchical. The project at Bolt Beranek and Newman (Woods and Makhoul 1973) demonstrates this approach. It involves a number of sub-components working together without a strict chain of command, i.e. any one can pass information to any other one directly whenever some particular knowledge is needed.

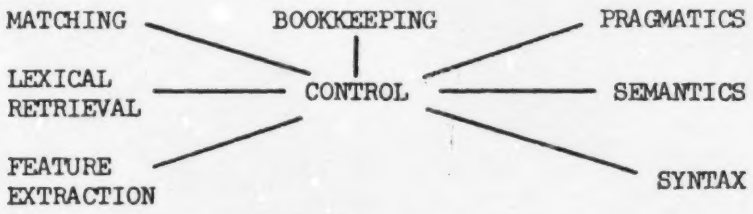

Heterarchical system organization (Winograd 1974)

Processing might proceed as follows (Winograd 1974):

> The feature extractor looks at the incoming wave forms, and suggests possible phonetic features. These can be used by the lexical retrieval component to see what words are possibly there. The matching component takes a possible word and checks it against a piece of wave form to check the likelihood that it is actually there. The syntax component evaluates possible sequences of words...and a control box to decide what component will do what when.

The second system organization he calls pandemonium. It characterizes the work being done at Carnegie-Mellon University (Reddy et al, 1973):

Pandemonium organizational model (Winograd 1974)

He describes this approach:

> It is as if we had a group of experts working on a common task, but no one of them knew anything about the others. Each expert might not even know how many others there were, or what kind of things they dealt with. Communication is managed by having each expert know how to propose sequences of words and assign a degree of confidence to them. Any individual component can look

at a current set of hypotheses (in the SCRATCHPAD), and either add new ones to it, or change the level of confidence in one of the old ones....

This domain certainly requires much cooperation between linguists and A.I. researchers if the effort is to achieve any kind of success.

6. Conclusions

While A.I. researchers have continually borrowed from linguistics, the converse has been rare. In their borrowings they have tended to ignore linguistic controversies and the unsettled state of theories. But they have also revealed some of the inadequacies, evasions, and errors in these theories by the practical necessity of programming a computer to implement them. If there are differences in opinion with respect to methodological principles, these should be overshadowed by the overriding similarity in goals--namely, to understand natural language.

NOTES

1. Work on machine translation using digital computers actually began in the late 1940's, but I am mainly concerned with the post-1960 period.
2. It is important to comment that models in A.I. need say nothing about human behaviour. Very often the concern with human activities is mainly as a source for ideas. There is no reason to expect that appropriate problem-solving techniques for computers will resemble in any way comparable human behaviour.

ACKNOWLEDGEMENTS

This research was supported in part by the National Research Council of Canada under grant No. A-5552. The hospitality of the Department of Electrical Engineering and Computer Science of the University of California at Berkeley for the fall and winter quarters of 1974-75 is warmly appreciated.

REFERENCES

ALPAC (1966) (Automated Language Processing Advisory Committee). Language and machines - computers in translation and linguistics. Nat.Acad.Sci., Washington, D.C.

Bar-Hillel, Y. (1960) The present status of automatic translation of languages. In Advances in Computers, Vol.1, F.Alt (Ed.) Academic Press, N.Y., pp. 91-163.

Bobrow,D.G. (1963) Syntactic analysis of English by computer - a survey. Proc. Fall Joint Comp. Conf.24 Spartan Books, Baltimore, pp. 365-387.

_____ (1964) Natural language input for a computer problem-solving system. In M.Minsky (1968).

Celce-Murcia,M. (1972) Paradigms for sentence recognition. U.C.L.A. Dept. of Linguistics.

Chomsky,N. (1957) Syntactic Structures, Mouton and Co., The Hague.

_____ _____ (1965) Aspects of the Theory of Syntax, The M.I.T. Press, Cambridge.

Coles,L.Stephen (1972) Techniques for information retrieval using an inferential question-answering system with natural-language input. Technical Note 74, SRI Project 8696, Stanford Research Institute, Menlo Park, California.

Dreyfus,H.D. (1972) What Computers Can't Do: A Critique of Artificial Reason Harper and Row, N.Y.

Feigenbaum,E. and Feldman,J. (1963) Computers and Thought McGraw-Hill, N.Y.

Fillmore,C.J. (1968) The Case for Case. In Bach and Harms (Eds.) Universals in Linguistic Theory Holt,Rinehart and Winston, N.Y., pp. 1-90.

_____ (1971) Some problems for case grammar. In O'Brien (Ed.) Linguistics: Developments of the Sixties - Viewpoints for the Seventies 22nd Annual Georgetown Roundtable. Monograph Series on Languages and Linguistics, No.24, Georgetown pp. 35-56.

Green,B.F.,Jr., Wolf,A.K., Chomsky,C., and Laughery,K. (1961) Baseball: an automatic question answerer. In Feigenbaum and Feldman (1963) pp. 207-216.

Green,C.Cordell, and Raphael,B. (1968) Research on intelligent question answering systems. Proc.ACM 23rd National Conference. Brandon Systems Press, Princeton, N.J. pp. 169-181.

Halliday,M.A.K. (1970) Language structure and language function. In J.Lyons (Ed.) New Horizons in Linguistics Penguin Books, England, pp. 140-165.

Harris,Z. (1960) Project summary in Current Research and Development in Scientific Documentation, No.6, Nat.Sci.Foundation, pp.52-53.

Kintsch,W. (1972) Notes on the semantic theory of memory. In E. Tulving and W.Donaldson (Eds.) Organization of Memory Academic Press, N.Y. pp. 249-309.

Kuno,S. (1967) Computer analysis of natural languages. In Proceedings of Symposia in Applied Mathematics Vol.XIX Am.Math.Assoc.pp.52-110.

Lakoff,G. (1970) Linguistics and natural logic. Synthese 22 pp.151-271.

Lindsay,R.K. (1963) Inferential memory as the basis of machines which understand natural language. In Feigenbaum and Feldman (1963) pp. 217-233.

McCarthy,J.T. (1958) Programs with common sense. In Minsky (1968) pp. 403-417.

McCawley, James T. (1968) The Role of Semantics in a Grammar. In Bach and Harms (Eds.) op.cit., pp. 124-169.

_____ (1969) Meaning and the description of languages. In C.Travis and J.F.Rosenberg (Eds.) Readings in the Philosophy of Language, Prentice-Hall (1971) pp. 99-129.

_____ (1972) A program for logic. In D.Davidson and G. Harman (Eds.) Semantics of Language D Reidel,Dordrecht p.498-544.

Minsky,M.(ed.) (1968) Semantic Information Processing M.I.T. Press, Cambridge.

Newell,A. et al (1973) Speech Understanding Systems: Final Report of a Study Group North Holland/American Elsevier, N.Y.

Norman,D.A. (1973) Memory, knowledge, and the answering of questions. In R.L.Solso (ed.) Contemporary Issues in Cognitive Psychology: The Loyola Symposium V.H.Winston and Sons, Washington, pp. 135-165.

Raphael,B. (1964) SIR: Semantic information retrieval. In Minsky (1968) pp. 33-145.

Reddy,R.D. et al (1973) The Hearsay speech system: an example of the recognition process. In Advance papers, Third International Joint Conference on Artificial Intelligence (3rd IJCAI) pp. 185-193.

Robinson,J.A. (1965) A machine-oriented logic based on the resolution principle. J.ACM 12 pp. 23-41.

Rumelhart,D.E., Lindsay,P.H. and Norman,D.A. (1972) A process of long-term memory. In E.Tulving and W.Donaldson op.cit. pp. 198-248.

Rustin,R. (Ed.) (1973) Natural Language Processing Algorithmics Press, N.Y.

Sandewall,E. (1971) Representing natural language information in predicate calculus. In B.Meltzer and D.Michie (Eds.) Machine Intelligence 6. American Elsevier,N.Y. pp. 255-277.

Schank,R. (1972) Cognitive dependency: A theory of natural language understanding. Cog.Psychol 3,4 pp.552-631.

_____ (1973a) The conceptual analysis of natural language. In Rustin (1973) pp. 291-309.

_____ (1973b) Identification of conceptualizations underlying natural language. In Schank and Colby (1973), pp. 187-247.

_____ and Colby,K.M. (1973) Computer Models of Thought and Language W.H.Freeman and Co., San Francisco.

Simmons,R.F. (1965) Answering English questions by computer: a survey. Comm.ACM 8,1 pp. 53-70.

_____ (1970) Natural language question-answering systems: 1969. Comm.ACM 13,1 pp. 15-30.

_____ (1973) Semantic networks: their computation and use for understanding English sentences. In Schank and Colby (1973) pp. 63-113.

Walker,D.E. (1973a) Automated language processing. Artificial Intelligence Center, Technical Note 77. Stanford Research Institute. Also in Carlos Cuadra (ed.) Annual Review of Information Science and Technology Vol.8, American Society for Information Processing, 1973.

_____ (1973b) Speech understanding through syntactic and semantic analysis. Advance papers 3rd IJCAI pp. 208-215.

_____ (1974) The SRI speech understanding system. Proc. IEEE Speech Symposium. Pittsburgh.

Weizenbaum,J. (1966) ELIZA - a computer program for the study of natural language communication between man and machine. Comm. ACM 9,1. pp. 36-45.

Wilks,Y. (1972) Lakoff on linguistics and natural logic. Stanford A.I.Project Memo AIM-170, Stanford University.

_____ (1973a) The Stanford machine translation project. In Rustin (1973) pp. 243-290.

_____ (1973b) An artificial intelligence approach to machine translation. In Schank and Colby (1973) pp. 114-151.

Winograd, T. (1971) Procedures as a representation for data in a computer program for understanding natural language MAC-TR-84 Cambridge, Mass: M.I.T.Proj.MAC, Ph.D.Thesis.

_____ (1972) Understanding Natural Language Academic Press, N.Y.

Winograd,T. (1973) A procedural model of language understanding. In Schank and Colby (1973) pp. 152-186.

_____ (1974) Five lectures on Artificial Intelligence. Stanford A.I.Laboratory Memo AIM-246, Stanford University.

Woods,W.A. (1968) Procedural semantics for a question-answering machine. Proc.AFIPS 1968 Fall Joint Comp.Conf. Vol.33, Pt.1, MDI Publications, Wayne,Pa. pp. 457-471.

_____ (1970) Transition network grammars for natural language analysis. Comm.ACM 13,10 pp. 591-606.

_____ (1973) An experimental parsing system for transition network grammars. In Rustin (1973) pp. 111-154.

_____ and Kaplan,R.M. (1971) The Lunar Sciences Natural Language Information System. BBN Report No.2265, Bolt Beranek and Newman, Inc. Cambridge, Mass.

_____ and Makhoul,J. (1973) Mechanical inference problems in continuous speech understanding. Advance papers 3rd IJCAI pp. 200-207.